

General SAM Hints and Tips

Adam Lyon (FNAL/CD/D0)
GCAS Meeting 12/19/02

Removing Special Runs
Cutting on Detector Quality
DB Query Tips
UDP problem
What to do if slow response

Checking Sam's Health
SamRoot
Future tools

Sam on clued0

Notes

- ◆ I'm not a super-SAM expert
- ◆ Have little experience with CAB
(see Heidi's talk)
- ◆ But, I'm involved in tools development with SAM (so I go to the SAM meetings)
- ◆ I know what the SAM team is worried about
- ◆ Heidi will give some tips too (Parallel jobs, Resubmission)

Removing special runs from dataset

- ◆ Problem: You want to make a dataset definition but exclude special runs

- ◆ Solution:
Use `onl_run_type` dimension.

`--dim='run_number 166780-166789 and
data_tier raw and onl_run_type Physics'`

Valid values are `Calibration`, `Cosmics`, `Physics`, `Special`, `Test`.
Note dimension queries are *case-sensitive* (use the capitalization above)

- ◆ Confusion: There is also a `run_type` dimension.
Valid values are: (`calibration`, `detector`, `detector data`,
`monte carlo`, `online archive`, `other`, `physics data taking`,
`run1`, `vlpc characterization`).
Use to separate collider data from MC. Note all lowercase.

Cutting on detector quality

- ◆ Problem: You want runs where the muon system had good quality
- ◆ Solution: Use `run_qual_group` and `run_quality` dimensions together.

```
--dim=' run_number 166780-166789 and data_tier raw and  
(run_qual_group MUO and run_quality REASONABLE)'
```

- ◆ Unfortunately, you can't do multiple groups (yet).

DB query tips

- ◆ Avoid querying on the file name when you can use meta-data (use `data_tier`, `physical_datastream_name`, `appl_name`, `run_number`, ...)

- ◆ If you must query on the file name, **never** put % first (e.g. %168573%)
 - ❖ You lose all indexing – DB will have to scan EVERY entry (millions).
 - ❖ Will be horribly slow and you will bog down DB server
 - ❖ If you find you can't avoid this usage, send mail to sam-admin so experts can come up with changes to the meta-data.

UDP problem (CAB/CLUED0)

- ◆ Jobs crash on CAB – a file already consumed appears to be delivered again
- ◆ Some SAM communication is done with UDP packets.
 - ❖ When a file is delivered, your project sends a UDP acknowledge packet to the station to say it got the file.
 - ❖ For some reason, station never receives packet.
 - ❖ Station assumes file wasn't delivered and it tries to deliver again.
 - ❖ But your job in fact consumed the file. When it is delivered again, your job throws an exception which is not caught by the framework.

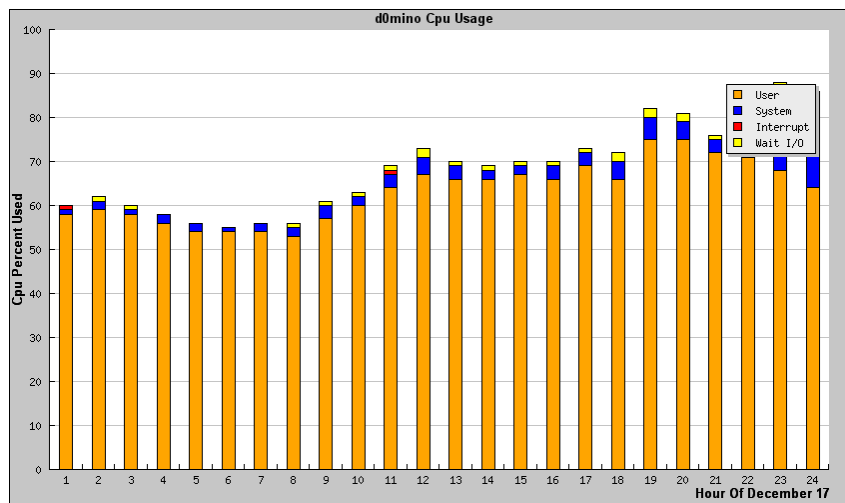
UDP problem

- ◆ Seen mostly on CAB
 - ❖ Seems to be correlated with CPU load
 - ❖ Increased timeout on station (d0mino station already upgraded). Stopped crashes.
- ◆ Seen rarely on CLUED0
 - ❖ Load is much lighter
- ◆ Long term solution is to move to CORBA for SAM communications
- ◆ Not seen on d0mino: all communication is internal

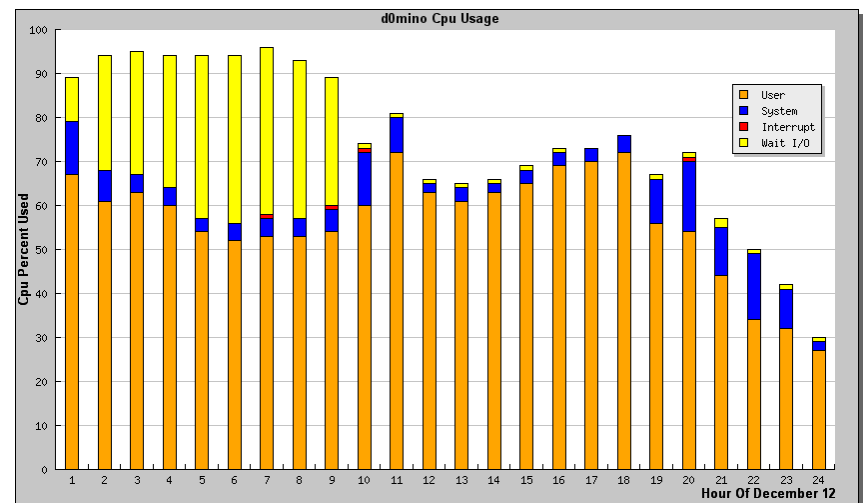
What to do when response is slow

- ◆ Problem: sam commands take a long time to respond
- ◆ This may be a d0mino system problem. **Check the d0mino performance metrics.**

<http://d0om.fnal.gov/d0admin/sysperf/>



A good day (12/17)



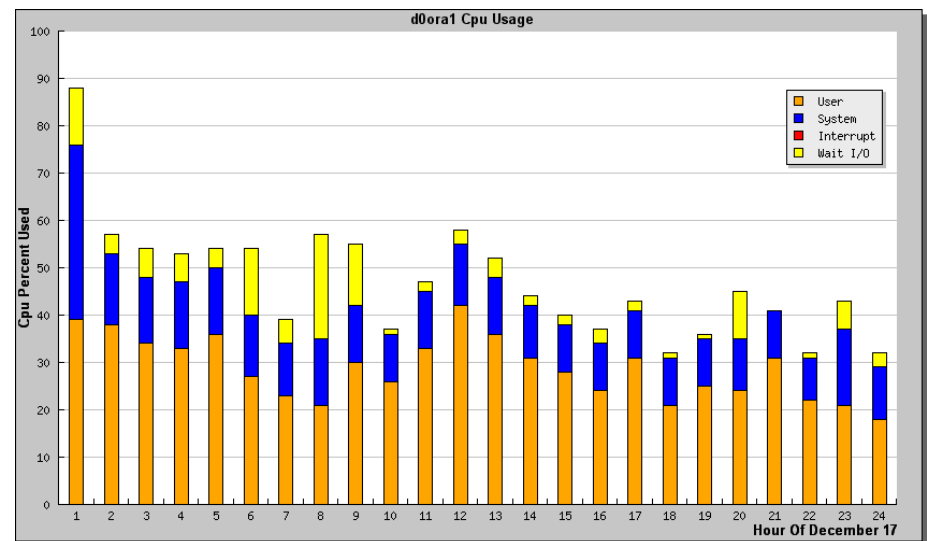
A very bad day (12/12)

A root job was writing to a full disk. Irix got confused and d0mino became bogged down.
(Root/Irix problem)

Slow response continued

- ◆ Problem: Misweb/DDE web sites are slow
- ◆ Usually caused by usage of d0ora1 (production oracle DB server). **Everybody** talks to d0ora1.

- ◆ Note that response will be slow early morning (FNAL time) due to backups and other jobs.



- ◆ If CPU is maxed then DB response will certainly be slow. But DB can be slow for other reasons too.

Is SAM healthy?

- ◆ A quick check is:

```
> setup sam  
> time sam locate foo  
No such file: foo  
1
```

- ◆ Timing on
clued0 should be around 2.5s user, 0.05s system
d0mino is longer (8s user, 5s system)
If the system time >> user time, something could be
wrong with the system

SamRoot

- ◆ Instead of copying root files to your area with SAM, run SAM within root (Sam-Root Serial Interface)
- ◆ See <http://d0db.fnal.gov/sam/doc/userdocs/SamRoot.html>
- ◆ Use *sam_client_api* A user interface to SAM from C++/root
- ◆ Caveats:
 - ❖ Only D0mino
 - ❖ Can only read files, not store
 - ❖ It's version may be (slightly) different than D0RunII
- ◆ Like runProject.py, but it runs within root!
- ◆ Are people using this interface (Sam-team has gotten little feedback)? Please give it a try. Send experiences (good & bad) to sam-admin.

Sam response to problems

- ◆ In the past, response by SAM shifters to problems (e-mail to sam-admin) has been less than stellar – SAM team is trying to improve situation
- ◆ Kin Yip is now SAM shift coordinator
- ◆ Kin and Wyatt track mail to sam-admin and responses
- ◆ Shifter training improving

Future user features

- ◆ Working on tools for SAM project archeology
- ◆ Error handling for `runProject.py`
 - ❖ Your job copies files and your disk fills up:
SAM thinks you got all requested files
 - ❖ For now, you should be keeping track of copy or other system command errors

Sam on Clued0

- ◆ Installing/Configuring/Testing SAM on clued0 has been a long process. Thanks to Sinisa, Lukas, Dugan, ...
- ◆ Clued0 officially opened for SAM business on December 6, 2002
 - ❖ 20 stager nodes (so max 20 SAM jobs)
 - ❖ SAM only available through batch system
 - ❖ Limited to 6 simultaneous downloads from enstore/d0mino (< 1TB/day)
 - ❖ No stagers in DAB due to network hubs
 - ❖ Clued0 SAM-admins are Adam Lyon & Jon Hays
 - ❖ If problems, send mail to sam-clued0@fnal.gov. We triage and will forward to sam-admin if we're stumped.

Sam on Clued0

- ◆ Turn on has been fairly smooth:
 - ❖ Stagers now automatically start when node reboots
 - ❖ Fcp daemon death halted deliveries on 12/12 – restarted and all was well
 - ❖ Usual batch system “features”
 - ❖ Make sure you change ‘central-analysis’ to ‘clued0’ for station name
 - ❖ Make sure you submit your job from a NFS mountable directory
 - ❖ Don’t specify a destination machine – the batch system will land your job on a SAM stager node
 - ❖ If you kill your batch job, please also
`sam stop project --project=yourName_12345`

Sam on Clued0 Stats

- ◆ From 12/6 – 12/18, 0.158TB (502 files) have been transferred to clued0 station
- ◆ No one has tried storing files
- ◆ We're still small potatoes

